RELEASE 1.0

12-99 ESTHER DYSON'S MONTHLY REPORT 15 DECEMBER 1999

META SERVICE PROVIDERS: THE INTERNET'S SS7 NETWORK

By Kevin Werbach

As the Net scales, its structure is changing. Though open protocols, abundant capacity and intelligence at the edges still represent the Net's fundamental architecture, some value is shifting to a new class of meta service providers -- neither at the edge nor in the core, but on a different plane. Basic transport and value-added features are becoming separated, with businesses growing up around distributed virtual networks. These meta networks offer enhanced performance, localization, anonymous communications and other benefits. Such platforms also enjoy positive feedback loops with network-hosted applications, contributing to their growth.

There's a model for these developments. The good old telephone network also has a separate distributed network-based system for delivering value-added services efficiently. It's called the Signaling System 7 (SS7) network. The analogy isn't exact, but thinking about Internet infrastructure developments in the context of SS7 leads to some important conclusions about how service providers, Websites and enterprises are likely to interact in the future.

Networks aren't linear (that's what makes them networks), and they don't grow and develop in a linear manner. The Net's remarkable commercial success in recent years has altered the path of its evolution. Today the Net faces far greater burdens than its creators ever envisioned... or could have possibly designed for. The hardware and transmission links that make up the Internet must constantly be upgraded. Throwing bandwidth and processing power at them isn't enough.

Eighteen months ago, when we last examined developments in the Internet backbone, we found a number of innovative approaches including quality of service (QOS) management, caching, virtual private peering and bandwidth brokering, all supporting the Net's continued scaling (see **Release** 1.0, 6-98). All these ===>

REGISTER NOW FOR PC FORUM
MARCH 12-15, AT WWW.EDVENTURE.COM!
EARLY REGISTRATION ENDS JANUARY 7.

INSIDE

THE INTERNET'S SS7 NETWORK 1
BETTA GET META 2
  SS7 history.
  Inktomi.
CONTENT DISTRIBUTION 6
  Akamai.
  Adero.
  Digital Island.
  Cable & Wireless.
DISTRIBUTED CENTRAL AUTHORITIES 8
  By Esther Dyson.
SERVICE PROVIDER SOLUTIONS 12
  Edgix.
OTHER FUNCTIONS 14
  Zero Knowledge.
  E-commerce applications.
CAN THE CENTER HOLD? 17
  The Net's operating systems.
RESOURCES AND CALENDAR 20
  Names, phones, dates, URLs.
PC FORUM 2000 INFORMATION 23

remain important. Some infrastructure challenges, however, cannot be solved at any one point in the network; they effectively require their own parallel meta-Internets. To build and deliver services over these virtual networks, meta service providers[1] are coming into being.

BETTA GET META

Meta service providers are the next step in the Internet's fragmentation of network services. Instead of vertically integrated carriers owning everything from the fiber in the ground to the services and customers riding the network, we currently have Internet service providers (ISPs) connecting customers, backbone providers routing traffic between ISPs, and hosting companies managing Websites and hosted applications. Now the very notion of delivering information is being disaggregated. Meta service providers offer value-added services to customers, who may be end-users, content providers or other service providers, using virtual end-to-end networks.

The basic concept of distributed service layers on top of raw Internet Protocol (IP) carriage is not new. The domain name system (DNS) that resolves alphanumeric addresses to machine-readable IP numbers has employed mirrored servers dispersed throughout the Internet since the early 1980s. Usenet news servers, dating back to the Arpanet days, are another example. For the first time, though, meta services are becoming substantial standalone businesses. Companies are deploying the infrastructure to deliver an unprecedented array of mission-critical functions through virtual networks.

These new layers improve performance, but they also shift the balance of power. As services depend more and more on distributed delivery platforms, those platforms will become the value core of the Net. To understand how, it's worth examining the history of the telephone network and its own meta service layer, the SS7 signaling network.

SS7: an out-of-band experience

Technical descriptions of the Internet often begin by distinguishing circuit and packet switching. The public switched telephone network (PSTN) is circuit-switched: A dedicated end-to-end connection is established for each call.[2] The Internet, in contrast, uses packet switching, with messages divided into small packets that are routed independently over the seemingly most efficient path at a particular moment (see appendix in **Release 1.0**, 6-98).

But that's not the whole picture. There is packet switching in the PSTN as well, in the form of the SS7 network. SS7 is a protocol developed by Bellcore (now Telcordia), the research and development arm of the Baby Bell local phone companies; it's now a global International Telecommunications Union (ITU) standard. Telephone switches talk SS7 with one another over a proprietary packet-switched network overlaid on the voice infrastructure. Like the Internet, the SS7 network includes routers (called signal transfer points) and database-driven servers (called signal control points).

---

[1] *The term is our own; we're surprised that no one seems to have claimed it in some other context.*

[2] *Some PSTN voice traffic travels over packet-switched networks, for example on congested international lines, but this is the exception rather than the rule.*

When you make a phone call, the local switch at your phone company sets up the call and performs other features such as calling-card authentication or identification of the carrier for a toll-free call, by sending packet-based messages over the SS7 network.[3] Carriers also use SS7 to support features such as caller ID and call waiting. When carriers interconnect their networks they also link together SS7 networks, so that information can be passed back and forth seamlessly.

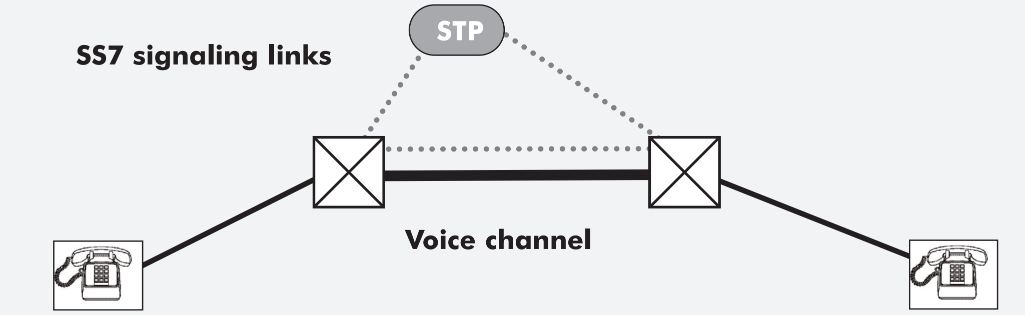

Figure 1 -- The SS7 signaling network in the telephone system.

SS7 demonstrated the benefits of separating call control and features from core service delivery. In the heart of the telephone network, it embodies the data networking model of discrete layers connected through standard interfaces taking the place of vertically integrated, proprietary solutions (see **Release 1.0**, 5-99).

The analogy between SS7 and the Internet's virtual service layers isn't perfect. The SS7 network uses physically separate links and different protocols from the telephone network's voice channels, while the Internet platforms all operate over the same basic IP infrastructure. Where SS7 handles only signaling, not the content of voice communication, the Internet services described below carry actual content.[4]

Finally, and perhaps most importantly, the Internet-based systems are decentralized and built on open (or potentially open) protocols.

---

[3] *The previous approach, in-band multi-frequency (MF) signaling, used special tones on the same channel that carried the voice communication. MF signaling was far slower than SS7, because end-to-end voice channels had to be opened just to carry the tones. Also, in-band signals couldn't be sent in the middle of a call, and couldn't be used with devices (such as database servers) that had no voice channel.*

[4] *A direct equivalent to the SS7 infrastructure would involve a parallel network carrying the border gateway protocol (BGP) messages that routers exchange to build routing tables. Such an architecture would be highly inefficient from a technical perspective.*

Christian Huitema, chief scientist in the Telcordia Internet Architecture Research Laboratory and formerly chairman of the Internet Architecture Board, explains: "This infrastructure is, in a sense, optional. The basic IP service remains available to whoever is satisfied with it, and can in particular be used to construct new applications." Then again, as we describe on page 17, openness isn't guaranteed in the future if the virtual service layers grow to become, in effect, the operating systems of the Net.

Inktomi: Just add scaling

The company closest to realizing an out-of-band distributed service layer for the Net is Inktomi. Inktomi first made its name as a provider of search-engine technology for HotBot, Yahoo! and other portals (see **Release 1.0**, 1-99), and then moved into caching with its Traffic Server product. Leveraging its core technologies of couple clustering, data flow and concept induction, Inktomi has built up two business areas: a suite of portal services products including search, comparison shopping and topic directories, which it sells on a private-label basis to portals and destination sites, and a set of network products built around Traffic Server that it sells to ISPs and other service providers.

Though Traffic Server began as a caching product, cto Paul Gauthier says that Inktomi's real goal for its network business is to add greater intelligence throughout the Net's backbones. "We're adding the pieces to the network to let information be manipulated at a much higher level," he explains. Switches and routers, even the so-called content-aware layer 4 devices, see only streams of packets passing through, but Inktomi's Traffic Servers store -- and therefore can recognize and act on -- entire Web pages, media streams, or other forms of content.

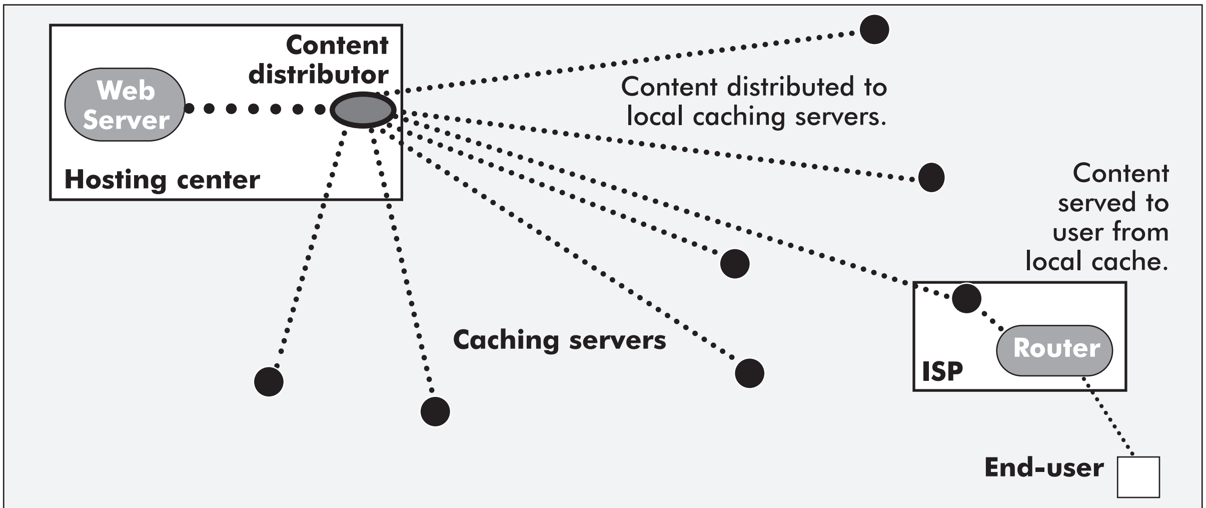
Content distributor
Web Server
Hosting center
Content distributed to local caching servers.
Content served to user from local cache.
Caching servers
ISP
Router
End-user

Figure 2 -- Distributed caching.

Gauthier explains that once Traffic Server caches are distributed through a network, they can also perform other functions such as filtering, reporting or transforming content to fit on a device such as a wireless phone. The smart caches can also be used for content distribution, by pushing out and synchronizing content to the edges of the network.

Inktomi uses the same “ingredient provider” business model for its network products that it successfully developed in the portal world. Most of the leading content-delivery services, including Adero, Sandpiper and Cable & Wireless’ Internet Shock Absorber, use Traffic Server as the basis for their offerings. Each company differentiates itself with the technology and networks it puts together to form a complete service offering. Virtually the only major player in the space that doesn’t use Traffic Server is Akamai, which has developed its own caching platform.

In September, Inktomi acquired Webspectives Software, a vendor of content-distribution and mirroring software. Using the Webspectives technology, now called Content Delivery Suite, Inktomi’s network of caches can communicate with one another across the network, cooperating to distribute, manage and deliver content most efficiently.

Gauthier agrees that Inktomi’s network of cooperating caches effectively becomes an out-of-band SS7-like signaling channel for the Net. “The phone network changed when people put computers into it. New services became possible,” he says. “We’re trying to sprinkle general-purpose powerful computers throughout the Internet infrastructure to create a whole new set of capabilities and services.”

Inktomi’s approach blurs the line between the edge and the core of the network. Rather than users at the endpoints requesting content from centrally located data-hosting centers, requiring constant transit of data back and forth, Inktomi uses intelligent nodes inside the network that push content close to where it’s needed.

The next step, Gauthier says, will be “content peering.” Just as backbone ISPs peer with one another to exchange traffic (see **Release 1.0**, 6-98), providers will negotiate agreements to exchange content more efficiently. As the technology vendor for many of the content-distribution networks, Inktomi may be able to develop a de facto standard protocol for connecting content-delivery networks and managing terms. Sandpiper and AOL, both Traffic Server clients, have already struck an agreement to push Sandpiper content directly into AOL caches.

The business models for content peering are yet to be worked out. For example, what level of granularity will companies use? Could a content-distribution network charge service providers more for receiving valuable *Wall Street Journal* content than for inexpensive AP wire stories, even though both require the same amount of bandwidth and storage capacity? Gauthier is optimistic an equilibrium will be reached.

For a broad standard to emerge, though, Inktomi and Akamai will have to agree...or one will have to win. Akamai rejects the concept of interconnection among content-distribution networks (see page 9). The two companies also diverge in business models (ingredient provider vs. turnkey service) and hardware platform (Sun vs. Intel). Not to mention West Coast vs. East Coast and UC Berkeley vs. MIT! With meta service layers becoming increasingly important, this will be a critical battle shaping the future of the Net.

CONTENT DISTRIBUTION: THE NEED FOR SPEED

Content distribution represents the first large-scale application of the meta service provider approach. Companies such as Akamai, Digital Island, Adero and Cable & Wireless have deployed thousands of servers and linked them into virtual networks to optimize delivery of content.

As we've described before (see **Release 1.0**, 6-98), the Net in its initial commercial form doesn't scale well enough because it depends on best-efforts delivery within and across multiple networks. Efficient statistical multiplexing, free interconnection between backbone networks and rapid increases in transmission bandwidth and router speeds made it possible for the Net to deliver acceptable performance in its early days. Going from thousands to millions to hundreds of millions of users on thousands of different networks is a manageable problem, so long as all the components of the Net are continually upgraded. The difficulty comes when a substantial number of those users want access to the same sites.

Yahoo! currently generates some 400 million pageviews per day, or more than 10 billion per month. Transporting 10 billion Web pages across the Web during the course of a month is one thing; delivering them from the same site is a far greater challenge. Busy sites such as Yahoo! must deploy clusters of hundreds of servers hosted at data centers operated by hosting providers such as Exodus and Frontier GlobalCenter, using sophisticated load-balancing algorithms to manage traffic across the server farms. Even so, information from these sites must traverse the Internet to reach customers, often facing congestion.

The challenges become more acute as increasing numbers of residential customers connect to the Net through broadband pipes. With over 1 million cable modem subscribers and 250,000 digital subscriber line (DSL) customers in the US alone, both growing rapidly, broadband to the home is finally becoming a reality. Customers with fast connections expect faster performance when requesting Web pages, which means that network throughput must improve.

Distributed intelligence

To meet this need, a group of startups is now selling content providers a new service called content distribution. Interest in the area has exploded ever since Akamai, the leading player, had a spectacular IPO in October. At this writing, Akamai has a market cap of $20 billion.

All of the content-delivery services use the same general architecture (see figure 2). When a user views a Web page under normal circumstances, a request is sent to a domain name server that routes it to the hosting Website, which then sends back the various page elements (text, images, executables, streaming media, etc.) to the user. Each element must traverse the network, often encountering congestion along the way at slow links or peering points.

For large sites, host server capacity is also a significant bottleneck, especially during traffic spikes, for example during a major news event or a popular live event. It may not be economical for a site to deploy enough servers to handle peak demand, especially if many of those servers will go idle most of the time.

The solution is to deploy a network of caching servers at strategic edge points on as many backbone networks as possible, so that a server is likely to be close to users requesting the information. The content distribution players differ in the number and capacity of their servers. Akamai has deployed 1700 servers on over 100 networks in 30 countries, using rack-mounted Pentium-class PCs to keep costs down. Digital Island, thanks to its acquisition of Sandpiper, has 1200 Unix-based servers in 21 countries, and it announced plans earlier this month to spend $150 million expanding to more than 5000 servers in the next three years. Cable & Wireless keeps its caching servers on its own backbone, currently at 84 nodes worldwide.

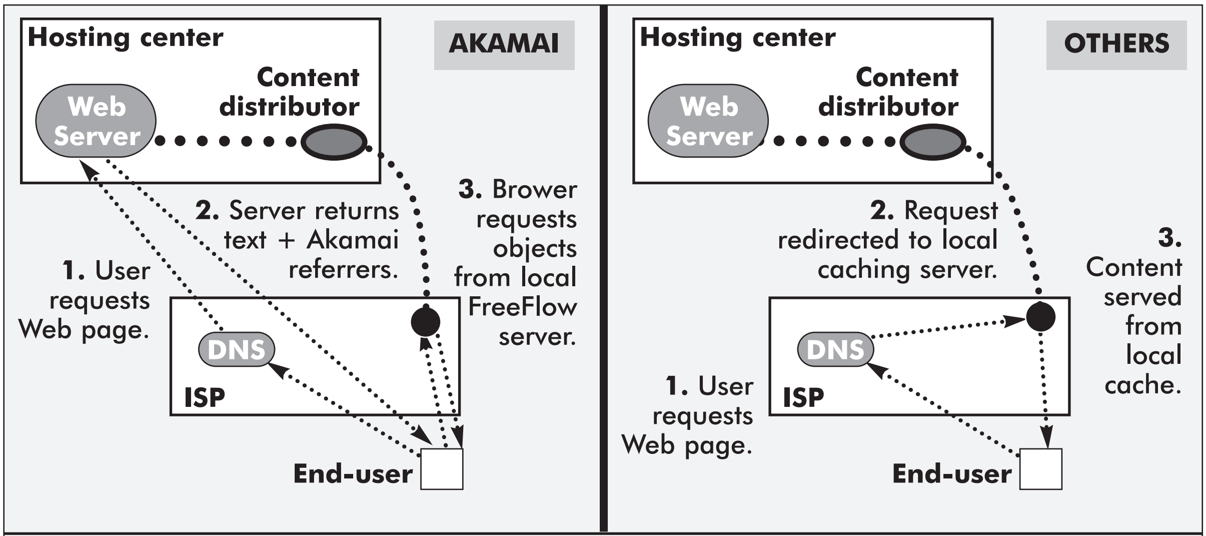
Hosting center
AKAMAI
Content distributor
Web Server
3. Brower requests objects from local FreeFlow server.
2. Server returns text + Akamai referrers.
1. User requests Web page.
DNS
ISP
End-user
Hosting center
OTHERS
Content distributor
Web Server
2. Request redirected to local caching server.
3. Content served from local cache.
1. User requests Web page.
DNS
ISP
End-user

Figure 3 -- Content distribution mechanisms.

When a user requests a page from a subscriber to a content-distribution service, his or her requests for larger components such as graphics are redirected to the caching server closest to that end-user. Akamai allows the DNS request to go through to the host server, where it is automatically redirected. Others such as Cable & Wireless and Adero redirect the IP addresses of customer sites, so that the end-user's DNS request resolves directly to the content-distribution network and doesn't hit the host site at all.

The technical challenge for these providers lies in the algorithms that efficiently direct traffic to the closest and most-available server in the network. Sometimes the first available server will be far away (physically or in terms of network topology) from the customer, or will be nearby but congested. Information must be kept synchronized so that users always get the latest version of a page. And the distribution function must be transparent to users and content providers. Early caching servers distorted usage counts and advertising statistics, but all the major content distribution networks now overcome this problem.

Akamai and Sandpiper (now merged with Digital Island) are the largest content-distribution providers. Adero and Cable & Wireless are more recent entrants; several other companies including Enron, Exodus, Mirror Image and Verio are jumping into the market as well using technology purchased from Inktomi (see page 4).

DISTRIBUTED CENTRAL AUTHORITIES

By Esther Dyson

One big issue facing the Net over the next few years is scaling. How can you have fine-grained control with millions of objects flying about? This isn't "control" as in regulation, but control as in managing a dispersed system according to central policies. It's not a system of a single central authority, but rather multiple competing authorities, each controlling its own objects and coordinating with other systems. Look at it as an immune system rather than as a digital nervous system; components know their own identities, follow the appropriate policies and recognize strangers.

As Kevin eloquently demonstrates, infrastructure isn't boring anymore. It will be especially not-boring as numerous e-commerce sites fail to deliver during the holiday rush, whether because of failing logistics or failing technology infrastructure.

But let's look longer-term and consider what new underlying technologies will be useful as we scale up. What will it take to manage the interactions of caches and filters and other meta systems when we have millions more things attached -- devices, resources and people?

In order for all this to work, the virtual world needs to be more self-aware, but with distributed intelligence. Each system may have its own central rules, but they often need to interact with those of other systems, and they need to be applied *locally* in local conditions. Those rules need to apply to and recognize everything from physical data routes, to what category of information is being transmitted, what applications are being used, what security restrictions apply. We need better categorization of all the objects flowing through the system and the services at the ends and middles. And we need a better way of representing rules that can be developed, if not centrally, at least by some controlling authority and applied locally. Agents need to be able to interact locally, but they need to follow explicit rules from the relevant authority.

You'll be hearing more and more about so-called policy-based "X" management, where X means everything from security to resource allocation and dynamic pricing, spam filtering, geography-based taxation and other regulations, data-handling (privacy) standards, translation services and traffic priority. That requires well-defined rules and policies and conflict resolution -- and standard definitions of what the rules apply to.

We'll need more standards, and more directories that manage roles, categories, permissions, ranks and the like, of which Novell's directory services are an excellent example. We'll also need application detection, somewhat like virus detection (but applications aren't usually trying to hide their nature; they're simply obscure because no one thought to identify them). This will require more cooperation between applications and between application developers. As Kevin points out, it may mean that Cisco, not Microsoft, controls the real operating system of the Internet.

Akamai: go with the FreeFlow

“The Internet is a giant theoretical network problem if I ever saw one,” says Akamai president and coo Paul Sagan, describing the company’s academic origins. Inspired by World Wide Web creator Tim Berners-Lee (whose office was down the hall), MIT applied mathematics professor Tom Leighton and graduate student Danny Lewin led a team that developed algorithms to optimize dynamic routing of Internet content across a distributed server network. They built a business plan around their ideas, and the rest is -- so far -- a classic Internet success story.

The hiring of Sagan (previously president of Time Inc. New Media), and of former BBN head George Conrades as ceo gave Akamai a high-profile management team, and the company quickly signed several top Web content providers as beta customers for its FreeFlow service. It received venture investments from Battery Ventures, Polaris Venture Partners, TCW and Baker Communications Fund, along with strategic investments from Apple, Microsoft and Cisco, before going public with a bang in October.

Akamai now has over 300 employees (including 40 PhDs) and over 100 customers ranging from behemoths such as Yahoo! and CNN to up-and-coming sites such as Listen.com. Customers run a software utility to tag objects for delivery through Akamai’s content-distribution network. Akamai charges them based on bandwidth, at $1995 per megabit per second.

Sagan says the challenges Akamai has overcome are harder than they may appear: “We don’t mirror sites, we replicate only the content that’s in demand, near where it’s demanded.” Akamai’s proprietary algorithms and architecture allow the complex calculations necessary to identify the most-efficient server without a central point of failure. “Any machine or node or region has the ability to make the calculation for any and every request,” says Sagan. When FreeFlow servers or entire networks go down, Akamai redirects requests to operational servers elsewhere.

The FreeFlow caching servers constantly send back data about networks loads and performance to Akamai’s network operations center, allowing Akamai to generate a real-time map of Internet traffic conditions, which it uses to determine the most efficient servers to route requests through. When the number of requests to a given site increases, Akamai’s algorithms automatically disperse content to more servers so that each one is less congested.

Sagan agrees content distribution is the first of many services Akamai can deliver: “It’s a real platform at the edges for applications.” This week, Akamai announced its EdgeAdvantage initiative to deliver application functionality such as profiling, log analysis, transaction processing, ad insertion and content transformation through its server network.

Akamai also announced the Internet Content Adaptation Protocol (ICAP) Forum, an open standards effort involving more than 30 companies that plans to submit a proposal to the Internet Engineering Task Force for communication between edge content devices and application servers. However, Sagan sees no value in competing content-distribution networks talking to one another. Akamai can address the problems facing its customers alone, he argues. Communication between meta networks, he says, could commoditize the service layer in favor of software and hardware vendors, but wouldn't benefit content-providers or end-users.

Adero: international content distribution

Adero ceo Jonathan Crane, former evp of national and multi-national accounts at MCI, knows a thing or two about SS7 and other elements of the telephone world. Crane joined Adero in October, around the time the company launched its AderoWorld content-distribution service. "The issue is finding different ways to assist the Internet to become scalable and to improve its performance," he explains. "It's not just continuing to boost bandwidth availability."

Adero, based in Boston, MA, currently has about 70 employees. Investors include Spectrum Equity, Fidelity Ventures, M/C Venture Partners and Osborn Capital.

Adero focuses on delivering content efficiently across geographic and network boundaries, such as between sites in the US and users elsewhere in the world. Currently the majority of the popular content on the Net is located in the US, but the user base and content elsewhere in the world are growing rapidly. International connections are expensive, and carriers from other countries typically must pay the full cost to connect their data networks to Internet backbone exchanges within the US. International links also tend to be more congested than domestic backbones in the US. The way to get acceptable performance in other countries, therefore, is to minimize the amount of content that has to traverse international circuits to reach customers.

Adero is building a global network of caching servers in 55 countries. Crane says Adero has the "ability to move content to be deployed worldwide, so as the local user comes on the Net he's not going to the original website to find what he wants; he's getting it served locally."

Adero mirrors Websites on its distributed network of caches, using intelligent routing algorithms in its GeoTraffic manager software to ensure content remains synchronized and fresh. Adero is partnering with local network and caching providers in other countries to expand its global network. As cto Joe Bai explains, "We provide the requisite services to tie together these far-flung caches and to make them a business instead of a cost-avoidance mechanism."

Crane sees Adero's platform enabling several other services for content providers beyond content distribution. For example, US-based sites looking to capture audiences in other countries could potentially use Adero to manage localization of their content, including language translation and application of appropriate taxation and other legal requirements in the destination country. Content providers can do this localization themselves, but a network-based service platform such as Adero's could allow such localization to happen on the fly based on a customer's location, without requiring expensive duplication of servers and content.

Crane believes that just as SS7 gave the telephone network "shared intelligence," emerging network-based services on the Net will allow a new intelligent network layer to emerge in order to improve speed and facilitate service delivery.

Digital Island/Sandpiper: content distribution + conduit

We wrote about Digital Island when it was primarily a backbone provider focused on efficient delivery of international traffic (see **Release 1.0**, 6-98). The company has since developed additional services including TraceWare, which maps IP addresses to geographic regions so that content and e-commerce providers can more efficiently localize their offerings to users around the world.

In October, Digital Island agreed to acquire Sandpiper, which offers content-distribution services, for $630 million. The merger brought together Digital Island's backbone and data centers with Sandpiper's network of intelligent content replication and distribution servers. Sandpiper's Footprint service, available since September 1998, includes servers on more than 20 different networks. Customers include Earthlink, CNBC.com, Blue Mountain Arts and PBS, and Sandpiper also handled content-distribution for several live events including the 1999 Academy Awards and the *Drew Carey Show's* live Webcam episode.

Mark Wagnon, director of network services for Digital Island, says that unlike competitors, "we're providing full processing services all the way out to the edge." Digital Island has deployed 1200 Sun Netra servers, capable of performing sophisticated transaction-processing functions in addition to caching static content. It has also placed switches and routers at its edge locations, making possible two-way functions such as secure e-commerce directly through its distributed overlay network. Because Digital Island has deployed a global managed IP network, including private peering relationships with the major backbones, it can offer reliable performance across its infrastructure for hosting, content distribution, and other services.

Wagnon thinks meta service layers will be particularly important for e-commerce as Internet appliances and other access devices such as set-top boxes become more prevalent. "You have to have a two-way network out near the edge to capture that impulse buy," he argues.

This month, Sandpiper announced a deal with Microsoft to integrate multicast streaming media into its content-delivery network. On a global basis, Sandpiper has more than six gigabits of streaming capacity, allowing it to support millions of simultaneous users. The company also announced a partnership with Inktomi and Sun to greatly expand the scope of its network.

Cable & Wireless: We got network!

One striking aspect of the content-delivery market is that the early players don't operate their own backbone networks. Because these are virtual services that seek to enhance performance across the entire Internet, there isn't necessarily an advantage to owning the underlying transmission links.

Cable & Wireless (C&W) believes otherwise. C&W has been offering domestic and international telephone service in Europe, the Caribbean and elsewhere for more than a century. It became a major Internet backbone player when it bought MCI's Internet network, which MCI divested in connection with its acquisition by Worldcom in 1998. However, C&W has so far largely failed to capitalize on this asset, as many key employees and customers deserted following the ownership change.

In the content-distribution arena, C&W feels it can leverage its network to offer a more reliable and lower-cost service than its competitors. Eliza Wojtaszek, product manager for C&W's Internet Shock Absorber (ISA) service, argues that because the company can monitor and maintain quality of service across its network, it can quickly resolve any network problems that interfere with the effectiveness of the content-distribution service. On the other hand, because C&W only hosts servers on its own network, it can't eliminate slowdowns on other networks between its backbone and end-users.

ISA uses "anycast" addressing to deliver requests to its caching servers. Customers remap the IP addresses of their sites to a new address supplied by C&W (see figure 3). That same address is assigned to all the nodes in C&W's content-distribution network. When a user requests a page, the local DNS server resolves to the ISA address. Through standard Internet routing mechanisms, the request automatically flows to the closest server matching that IP address. If the ISA server doesn't have the requested objects cached, it pulls them from the host server and serves them to the customer. Eliza Wojtaszek, C&W's product manager for the service, says this approach enhances reliability: "If one of the servers fails, other servers will pick up the load, just like the routers on the Internet."

C&W also hopes to compete on price, charging content providers for megabits of data transferred from its caches rather than the amount of bandwidth (in megabits per second) it makes available. "We actually charge for the content that is being served," says Wojtaszek. "Other competitors charge for the rate that content is being served." Because C&W has so much extra capacity in its network, she explains, it need not throttle the bandwidth its ISA customers have available.

Other backbone providers and carriers are bound to jump into the content-distribution game. Only the largest network operators have sufficient coverage to deliver this service entirely on their own infrastructure, but carriers and meta service providers are likely to develop partnerships to expand their offerings. Combinations involving Web-hosting providers and next-generation exchange points such as InterNAP (see **Release 1.0**, 6-98) and Equinix should also develop.

SERVICE PROVIDER SOLUTIONS

Content distribution helps get information from content providers to users throughout the network. At the other end of the process, those users depend on local ISPs for the quality of their online experience.

ISPs benefit from content-distribution services because they lead to better performance for their users, and reduce the ISPs' bandwidth requirements (which is why so many of them have been willing to host caching servers from companies such as Akamai and Sandpiper in their networks). But ISPs can't guarantee that their users will visit the same sites that subscribe to the content-distribution services. Instead, ISPs use caches to improve performance on the sites their users request most frequently.

There's some classic Internet business model uncertainty here. Both Websites and ISPs may carry the same content, but they pay different service providers to improve delivery speeds. Inktomi's notion of content peering (see page 5) could be the solution, but it doesn't yet exist in the real world, and others such as Akamai have different ideas.

Last year, we described some of the solutions being marketed to ISPs, including caches and virtual private peering (see **Release 1.0**, 6-98). A new player in this market is Edgix, which hopes to leapfrog its more established competitors.

Edgix: living on the edge

Edgix, based in New York City with a development team in Provo, UT, uses a combination of satellite and wireline networks to intelligently prepopulate ISP caches with popular information. The company presently has more than 35 employees, and funding from angel investors and Novell Ventures; it is currently raising an institutional funding round. The Novell connection exists on several levels: Edgix uses Novell's caching servers in its network, and its founders Kyle Powell and Drew Major and Mark Hurst were early Novell employees responsible for the creation of NetWare. The third founder, cto Mark Hurst, developed the text editor that eventually became WordPerfect.

Skycache (see **Release 1.0**, 6-98) and iBeam similarly employ satellites to distribute Internet to ISPs and avoid Internet congestion. However, Edgix president and ceo Rangu Salgame argues that there is little first-mover advantage. “None of the players have figured out a sustainable revenue model yet,” he points out. Salgame was president of Bell Atlantic Internetworking before taking the helm at Edgix. The company also hired Jupiter Communications broadband analyst Abhi Chaki to head up business development.

“We bring the content that's most relevant to the service provider's user population,” states Salgame, explaining that Edgix's proprietary heuristics use a mix of popularity and top-down criteria to send content out to caches near the edge of the network. By employing a mixture of delivery technologies, Edgix believes it can deliver content in the most economic and technically efficient manner.

Where services for content providers use large numbers of servers dispersed throughout the Internet, Edgix focuses on bypassing congested backbones and delivering data directly to its ISP customers at the network's edges. “We are betting on the edge of the network,” says Salgame. Like his competitors, Salgame sees the game for Edgix as more than just content distribution. Edgix is working on localization and and streaming media services to incorporate into its service.

The company also plans to experiment with pricing models. Salgame says Edgix plans to charge a fixed monthly fee to ISP customers, but it will also allow ISPs to pay based on the value they derive from signing up for the service.

Edgix has (undisclosed) beta customers in several regions of the world. It plans to launch the service publicly by early February.

OTHER FUNCTIONS FOR PARALLEL CHANNELS

The meta service provider model isn't limited to content distribution. Says Telcordia's Huitema, "I fully expect that other such infrastructures will be deployed over time, to support commerce, advertising, gaming, and many more applications." The content-distribution players themselves are likely to offer additional functions from their servers, but other companies are developing their own virtual service layers for specific functions. The best example is Zero Knowledge Systems, though there are also interesting developments in the e-commerce realm.

Zero Knowledge: It's freedom, baby!

Zero Knowledge Systems, based in Montreal, Canada, announced its Freedom service for anonymous Internet communications early this year, but it has only launched commercially in the past few weeks. Implementation of the company's ambitious effort took longer than expected, but Zero Knowledge president Austin Hill says getting it fool-proof -- and expert-proof -- was essential. "We had no idea how hard it was," he admits.

Zero Knowledge was founded in 1997 by ceo Hammie Hill and his sons Austin and Hamnett (who serves as coo). Cryptography expert Ian Goldberg is the company's chief scientist, and Zero Knowledge has engaged Berkeley professor Bruce Schneier, a leading authority on security technology, to review and critique its technical implementation.

Online privacy is an important and many-faceted issue (see **Release 1.0**, 4-98). Encryption tools based on public-key cryptography have been available for several years. Numerous anonymous remailers allow e-mail messages to be sent without identifying source headers. Though an active community of "cypherpunks" and commercial software companies develop these tools, and many people use them, they have two primary limitations. First, existing solutions are often cumbersome to use. Second, even when the communication between two points is encrypted, the sender still needs to trust the servers on both ends to keep information secret. Moreover, encrypting the content of a message doesn't prevent tracking its origin and destination, which may be just as sensitive.

There's also a tradeoff in existing systems between privacy and persistent identity. Using a pseudonym, Al can send an untraceable message to Ethel through an anonymous remailer. But if Ethel wants to reply to the pseudonym, the remailer operator generally must keep a list of real user addresses and their aliases. Even when the operator is committed to privacy, legal authorities may force disclosure of user information. Such a situation forced Julf Helsingius, who operated the leading anonymous remailer at anon.penet.fi, to shut down his service in 1996 when confronted with a subpoena in a lawsuit filed by the Church of Scientology.

For Web surfing, the lack of anonymity in the middle of the network is an even greater problem. Every packet has originating and destination IP address information in its headers, meaning that someone viewing a packet (or a Web server log) can identify a particular user visiting a site. Anonymizer services exist to scramble this information and other identifiers such as cookies, but users still need to trust them to maintain security. Moreover, anonymizers don't allow users to browse the Web with aliases. You can either be yourself or no one, but you can't create a consistent fictional personality (pseudonym) without allowing it to be linked back to you.

My own private Internet

Zero Knowledge solves these problems through client software to make management of anonymous and aliased communications simple, and by creating its own distributed anonymous service layer across the Net. The company has located about 150 Freedom servers at ISPs in 25 countries. ISPs participate because they want to promote privacy, and also for more mundane benefits: Zero Knowledge gives them a share of revenues from Freedom subscriptions based on the amount of bandwidth they contribute.

Like the content delivery services described earlier, Freedom involves a distinct virtual channel on top of the general-purpose Internet to enable new functionality. “This is a complete privacy infrastructure,” explains Hill. “We have a complete IP overlay network, with routing tables, public keys, digitally signed messages, key distribution and optimization of packets for throughput.” Zero Knowledge effectively rebuilt the Internet from the ground up, redesigning everything from routing tables to routine network statistical reporting to be secure and private.

When a customer sends a message or requests a Web page, the Freedom client-side software selects up to three servers and establishes a connection path through them to the destination. The user’s information is encoded in a set of nested cryptographic “envelopes” and sent to the first server. That server opens the first envelope and forwards the message to the next server in the chain. That first server knows the origin of the message but not its contents or destination, because that information remains encrypted. This process is repeated until the message reaches the final Freedom server, which knows the destination address but not the origin of the message. The information is then sent to the recipient, who can determine only the last Freedom server in the chain, not the origin of the message.

Zero Knowledge sells Freedom to customers for $49.95 per year including five “nyms” or pseudonymous identities. Hill says there is already a rush to claim the best nyms, much like the land grab for easy-to-remember dot com domain names. To ensure smooth performance as it builds out its network, the company will cap new customer signups at ten thousand per week until February.

Hill sees anonymous identities as the first of many services the Freedom infrastructure makes possible. Ultimately he expects the network to be outsourced to service providers, with many companies deploying privacy-enabled applications using open standards and software-development kits that Zero Knowledge is creating.

Meta services for e-commerce

Distributed virtual service layers can also be used for e-commerce. The most valuable information to online merchants is the metadata surrounding transactions, such as user preferences and purchasing histories (see **Release 1.0**, 9-98). Such data allow e-commerce sites to personalize and target their offerings, enhancing conversion rates and providing a better experience for customers.

Many tools exist to analyze user behavior as reflected in server log files, to offer recommendations through collaborative filtering or other mathematical techniques, to segment audiences, or to develop profiles to personalize content and merchandise displayed. These services function more effectively the more data that is fed into them. Sites can develop databases derived from their own traffic, but they only see a fraction of their customers' online activity.

The next step therefore, is to aggregate information about user behavior across a multitude of sites. Just as content-distribution services make it appear to end-users that the Net simply operates faster, networked e-commerce services could make it seem as though sites simply offer more relevant merchandise and familiarity with their customers.

Such aggregation must be done carefully to avoid inappropriate or unwelcome use of personal data, especially when the cultural and legal boundaries in this area are both uncertain and subject to local variation. Merchants can offer benefits to their customers if they know something about them. However, customers should be able to draw the line and place limits on use of their personal data, rather than having all their online behavior automatically become an open book.[5]

The anonymous profile databases companies such as Doubleclick and Engage employ to serve targeted banner ads are a starting point, but these companies generally combine the database lookup and the ad serving into the same process. In the meta service provider model, the distributed virtual service is offered on a separate channel.

There aren't yet any good examples of this kind of virtual e-commerce service, but several companies are moving in this direction. For example, Trivida, a startup based in Culver City, CA, is pushing the idea of networked personalization. Trivida offers sites the ability to personalize their interaction with users, without purchasing expensive server software, by modeling user behavior across all the sites in its network. Coming at the problem from a different angle, some middleware vendors such as Software Technologies Corporation are working to expand the functional abstraction layers they have built on top of enterprise networks to operate across the Internet.

Alexa, the service that anonymously tracks users' paths through the Web and uses the data to recommend related sites (see **Release 1.0**, 5-98), is also developing meta services for e-commerce. Alexa is now part of Amazon.com, so it's looking for synergies between its technology and its parent's business objectives. The company's initial answer is a new service called zBubbles, which pops up product recommendations (available through Amazon.com and elsewhere) based on Alexa's database of Web-wide user behavior.

---

[5] *One benefit of Zero Knowledge's service (see page 14) is that it allows users to interact with sites using persistent pseudonyms. A user could visit an online drugstore regularly as "Superman" and develop a profile, but that profile couldn't be mapped back to the real person behind the pseudonym. Users can exercise fine-grained control over how cookies (which identify a user to a site) are handled for each pseudonym.*

E-commerce meta services will also be valuable for addressing local regulatory and financial issues. If transactions take place at the edge of the network, they can be customized for country-specific currencies, promotions, taxes and distribution restrictions.[6]

CAN THE CENTER HOLD?

It's worth stepping back to consider what the phenomenon of meta service providers really means. Services such as content distribution and anonymous communication are valuable and viable businesses in their own right, but they are also part of something bigger.

The Net's operating systems

Meta service providers could become something akin to operating systems of the Net. In other words, the basic platforms upon which other services and applications are built. The Net already has low-level OS-like phenomena in the domain name and IP addressing systems, as well as Cisco's pervasive Internetwork Operating System (IOS) that runs the vast majority of the network's routers. These service layers make the network run, just as DOS and the BIOS handle low-level functions for Windows PCs. Meta service providers are creating something at a higher level of abstraction, comparable to Windows with its graphical user interface.[7]

Think of a Website such as Microsoft's MSN as an application, running on a computer called the Internet. The OS manages the application's inputs and outputs, allowing users to interact with it, just as content-delivery services transfer data efficiently between Websites and their audiences.

The OS also mediates the relationships between applications, and meta service providers are already beginning to step into this role. For example, Blue Martini (see **Release 1.0**, 3-99) just announced integration of Akamai's FreeFlow service with its e-commerce merchandising platform. Blue Martini customers can select Akamai's content distribution as a standard option when deploying their sites. Akamai has a similar deal with Vignette (see **Release 1.0**, 9-98) to integrate FreeFlow with that company's StoryServer product.

As we've described in the past, techniques such as syndication will allow aggregation of information on the Web according to user demand (see **Release 1.0**, 7/8-99). Meta service providers could provide the underlying transport in such a scenario, and could also bring together content with other services such as advertising, localization and trans-coding for non-PC devices.

---

[6] *Then again, meta service layers such as Zero Knowledge's Freedom network will allow users to effectively choose their jurisdiction, by making it appear that their requests originate from a particular country.*

[7] *This analogy shouldn't be taken too far. Unlike most PCs, the Net won't have just one OS. Though we expect many of the meta service layers to consolidate and interoperate, there's no fundamental need to have a single platform for all the different kinds of services that might be deployed in this manner. However, as we describe on page 18, there are some competitive issues to consider.*

Opening up the control points

It's impossible to talk about something as an operating system these days without considering the policy and competitive implications. Here we turn back to the SS7 analogy.

Bellcore, which developed SS7, also created a software layer to run on top of the SS7 network to control advanced features regardless of the underlying hardware of network switches. This broader platform, called the intelligent network (IN), was supposed to open up the telephone network to innovative new services. Despite grandiose promises, the success of IN has been limited. There are two main problems. IN assumes intelligence built into the center of the network. This approach made IN more cumbersome than nascent Internet-based services that used cheap standard components in the center of the network with intelligence at the edge.

The second limitation of IN is that, though it is designed to allow hardware from different vendors to interoperate, it's still a proprietary platform developed by the incumbent local telephone industry. Unlike the truly open technical protocols of the Net, IN didn't create an ecology of free experimentation and innovation.

In a series of proceedings in the 1980s and early 1990s known as the open network architecture (ONA) decisions, the FCC attempted to pry open the signaling and other advanced features being developed on the telephone network, so that competitors could develop and interconnect their own services. Under the ONA rules, a competitor such as a voice messaging or value-added network provider could petition local phone companies to make certain internal features and functions of their networks available as an "unbundled" service.

ONA became the model for the network unbundling requirements the US Congress wrote into law in the 1996 Telecommunications Act. Among other things, these rules have allowed digital subscriber line (DSL) wholesalers such as Covad, Rhythms NetConnections and Northpoint to use elements of established local phone networks to provide competitive high-speed data services.

On its own terms, though, ONA was a failure. Phone companies complained that implementation of the unbundling requirements was cumbersome and expensive, while potential competitors claimed the phone companies dragged their feet and prevented them from gaining access to the features they needed. Long, complex FCC proceedings produced detailed rules and technical mandates, but few companies were actually able to use ONA to develop competitive new services.

Fast forward to today, in the golden age of the Internet. There's no need for a Net version of ONA because the Net is inherently open, or so we've been led to believe. Openness is supposedly baked into the core of the network because all the intelligence and programmability lies at the edges. It doesn't matter what techniques router and switch vendors use to forward packets more quickly, since services are defined by end-users.

In reality, as we argued earlier this year with regard to broadband access services, the Net's openness is far from guaranteed (see **Release 1.0**, 2-99). The very advances that make possible new services and enhancements such as broadband connectivity also create possibilities for new bottlenecks. At the core of the network, there are no providers with monopoly control over an essential resource analogous to the cable industry's coaxial wires into customers' homes. However, if the meta service platforms now being deployed grow into something like operating systems for the Net, expect all the familiar antitrust ramifications.

Powerful meta service providers wouldn't necessarily be a bad thing. So long as the new Net OS providers recognize that they will only shoot themselves in the foot by acting in an exclusionary manner, having de facto standard virtual service layers throughout the Net is probably the only way Internet performance and features can keep up with demand. One big question is whether and how content peering will develop (see page 5) to allow distributed networks to exchange information.

Most of these services are just getting off the ground, and others have yet to be deployed, so it's premature to predict the competitive landscape that will emerge as they become more significant. But throughout the Internet's history, subtle shifts in infrastructure have had significant impacts on the success of major players, and the development of meta service providers should be no different.

Not so stupid after all?

In David Isenberg's memorable phrase, the Net is a "stupid network," which uses intelligence at the edges to take the place of the expensive, proprietary switches that dominate the telephone network. This notion and the broader concept of "end-to-end" system design (see **Release 1.0**, 5-99), help to explain how the Net can foster so much collective innovation while still remaining manageable and accessible to all.

Meta service providers are essentially putting intelligence back into the network, but it's a different kind of intelligence than the telephone network employed. The PSTN had one (proprietary) voice network and one (proprietary) signaling network to control it. The Net has many redundant networks for both information payloads and meta services, allowing different approaches to compete in the open marketplace. Where the PSTN ossified with expensive, complex switches, the Net is developing internal intelligence based on low-cost, off-the-shelf hardware running software that interfaces through open protocols.

Going forward, stupid and smart approaches will no longer be so diametrically opposed, because the lines between the edge and the core of the Net will no longer be so clear. As the network and its functions are increasingly defined by software, virtual abstraction layers will help users solve problems without worrying about the location of content or hardware. For some applications there may be value in tuning the network a particular way, but because the tuning takes place above infrastructure layer, it will be compatible with other uses.

It's not at all certain how smooth or rapid this evolution will be. There will be snags and unexpected developments, new winners and new losers. The one certainty is that the Net can't stand still, because the demands placed on it continue to grow.

RESOURCES & PHONE NUMBERS

Jonathan Crane, Joe Bai, **Adero**, (617) 241-3600; fax, (617) 241-5166; jcrane@adero.com, jbai@adero.com
Danny Lewin, Tom Leighton, **Akamai**, (617) 250-3000, fax, (617) 250-3001
Paul Sagan, **Akamai**, (617) 250-3000, fax, (617) 250-3001; paul@akamai.com
Eliza Wojtaszek, **Cable & Wireless**, (703) 287-6304; eliza.wojtaszek@cwusa.com
Mark Wagnon, **Digital Island**, (415) 738-4583; mwagnon@digitalisland.net
Abhi Chaki, **Edgix**, (212) 527-7527; fax, (212) 527-7528; abhi@edgix.com
Rangnath (Rangu) Salgame, **Edgix**, (212) 944-9662, rangu@edgix.com
Paul Gauthier, **Inktomi**, (650) 653-2800; (650) 653-2801; paul@inktomi.com
Eric Brewer, **Inktomi**, (510) 883-7330; brewer@cs.berkeley.edu
Leo Spiegel, **Sandpiper**, (805) 371-9098; leo@sandpiper.net
Christian Huitema, **Telcordia**, (973) 829-4266; fax, (973) 829-2504; huitema@research.telcordia.com
Bob Lucky, **Telcordia**, (732) 758-2100; rlucky@telcordia.com,
Austin Hill, **Zero Knowledge Systems**, (514) 286-2636; fax, (514) 286-2755; austin@zks.net

COMING SOON

- *How big companies innovate.*
- *The Net in the educational process.*
- *And much more... (If you know of any good examples of the categories listed above, please let us know.)*

Release 1.0 is published monthly except for a combined July/August issue by EDventure Holdings Inc., 104 Fifth Avenue, New York, NY 10011-6901; (212) 924-8800; fax, (212) 924-0240; www.edventure.com. It covers software, the Internet, e-commerce, convergence, online services, messaging, data networking, groupware, streaming media, enterprise applications, wireless communications, intellectual property and other unpredictable topics. Editor: Esther Dyson (edyson@edventure.com); publisher: Daphne Kis (daphne@edventure.com); managing editor: Kevin Werbach (kevin@edventure.com); office manager: Helen Martin (helen@edventure.com); marketing manager: Joanna Douglas (joanna@edventure.com); systems director: Scott Giering (scott@edventure.com); circulation assistant: Natasha Felshman (natasha@edventure.com); executive assistant: Bill Demeritt (bill@edventure.com); assistant: Kara Holmstrom (kara@edventure.com). Copyright 1999, EDventure Holdings Inc. All rights reserved. No material in this publication may be reproduced without written permission; however, we gladly arrange for reprints or bulk purchases. Subscriptions cost $795 per year, $850 overseas.

RELEASE 1.0 CALENDAR

2000

January 4-8	**Macworld Expo** - San Francisco, CA. What will Steve Jobs pull out of his hat next? Call (781) 551-9800; www.macworldexpo.com.
January 6-9	**Consumer Electronic Show** - Las Vegas, NV. Gizmos! We got gizmos! Check out the latest and greatest. Fax (805) 644-5793; www.CESweb.org.
January 16-20	**RSA Conference 2000** - San Jose, CA. New developments in data security and cryptography. Call (415) 544-9300; fax, (415) 544-9306; www.rsasecurity.com/rsa2000.
January 25-27	**Showcase** - Palm Springs, CA. Gina Smith and the editors of Upside hand-pick the top 125 startups of the past three months. Call (888) 338-7743; fax, (415) 908-2660; www.upsideevents.com/showcase/.
January 26-28	**The Fifth International Symposium on Artificial Life and Robotics** - Oita, Japan. New technologies based on simulation and hardware for the 21st century. Call +81 (97) 554-7831; fax, +81 (97) 554-7818; arobsecr@cc.oita-u.ac.jp, www.arob.cc.oita-u.ac.jp.
January 26-28	**Online Advertising 2000** - New Orleans, LA. Find out where all those VC-supported marketing dollars are going. To register, (800) 647-7600; fax, (800) 717-3237; info@worldrg.com; www.worldrg.com.
February 1-4	**LinuxWorld Conference & Expo** - New York, NY. For believers, heretics and others. Call (800) 657-1474; LinuxWorldexpo@idg.com; www.linuxworldexpo.com.
February 6-9	**Demo 2000** - Indian Wells, CA. Chris Shipley picks the hot startups. (650) 312-0545; fax, (650) 286-2750; www.demo.com.
February 14-15	**Milia 2000** - Cannes, France. Where the business of convergence takes shape. Contact Janice Johnston-Howie, +33 (1) 4190-4471; www.milia.com.
February 23-26	**TED 2000** - Monterey, CA. Richard Saul Wurman's annual uncategorizable gathering. For more info, call (401) 848-2299; fax, (401) 848-2599; www.ted.com.
Feb 28 - March 2	**Embedded Systems Conference Spring** - Chicago, IL. Tools, techniques and services for embedded system design. To register, call (888) 239-5563; esc@mfi.com; www.embedded.com/spring.
Feb 29 - March 2	**Internet and Electronic Commerce Conference & Exposition** - New York, NY. From business to e-business. For more info, call (218) 723-9130; fax, (218) 723-9122; www.iec-expo.com.
March 1-3	**Jupiter Consumer Online Forum** - New York, NY. Jupiter's flagship event. Call (212) 780-6060 x6424; fax, (212) 780-6075; www.jup.com/events/forums/forum2000.html.
March 7-10	***ICANN Open Meeting** - Location TBA, most likely in Africa. Hear about developments in domain name policy. For info, see www.icann.org.

RELEASE 1.0 CALENDAR

March 12-15	****PC Forum** - Scottsdale, AZ. Sponsored by EDventure Holdings. You read the newsletter; now attend the conference! Call Daphne Kis, (212) 924-8800; fax, (212) 924-0240; daphne@edventure.com; www.edventure.com.
April 3-6	**The Evolution of Language** - Paris, France. Third in a series on the evolutionary emergence of speech and language. For more info, evolang@infres.enst.fr; www.infres.enst.fr/confs/evolang/.
April 4-7	**CFP 2000** - Toronto, Canada. 10th anniversary event, focusing on policy and technical solutions for freedom and privacy on the Net. For info, call (202) 637-9800; fax, (202) 637-0968; www.cfp2000.org.
April 9-12	**AIIM 2000:** Technologies Driving E-Business - New York, NY. The world's largest event for business information management solutions. Contact Debbie Chapman, (301) 587-8202; www.aiim2000.com.
April 17-18	**DIMACS Workshop on the Management of Digital IP** - Rutgers, NJ. Technical, legal and business aspects of digital intellectual property management. Call (650) 725-3897; fax, (650) 725-4671; dabo@cs.stanford.edu; crypto.stanford.edu/DIMACS.

** Events Esther plans to attend.*
# Events Kevin plans to attend.

Lack of a symbol is no indication of lack of merit. The full, current calendar is available on our Website, www.edventure.com.

Please contact Joanna Douglas (joanna@edventure.com) to let us know about other events we should include.

PC Forum 2000

Convergence 2000: Making the virtual real, and the real virtual

We are pleased to announce the twenty-third annual PC Forum, to be held this year in Scottsdale, Arizona from March 12th to 15th.

> *Over the next few years, we're going to be settling cyberspace, making it more hospitable to normal folks. That means increased maturity -- both in what we (the industry) offer, and in consumers' familiarity with online capabilities and semiotics. The challenge for vendors is figuring out what consumers want, and understanding that they don't all want the same thing: Some want help; others just want the facts, thanks.*
>
> *Accordingly, we need to build up the social, legal, technical and information infrastructure that will make the virtual world feel more real. Vendors will call the result stickiness; would-be market entrants will consider it friction that keeps them out. Consumers will know only that it's easier to find their way around and to get what they want. The whole world will be virtually close and visible to us; the challenge will be to decide what we want to be connected to and when. It's an exciting world these days, and we invite you to come design it with us!*

Confirmed speakers include:

- **Steve Ballmer,** Microsoft
- **George Bell**, Excite@Home
- **Richard Bressler**, Time Warner
- **George Conrades**, Akamai Technologies
- **Jim Crowe**, Level 3 Communications
- **Shelly Lazarus**, Ogilvy & Mather
- **Maureen McDaniels**, Women.com
- **Carl Malamud**, Invisible Worlds
- **Stephen Pelletier**, Offroad Capital
- **Tony Salter**, Boxman
- **Marjorie Scardino**, Pearson
- **Eric Schmidt**, Novell
- **Mark Walsh**, VerticalNet

Plus several other speakers and company presenters to be announced!

To register, visit our Website at **www.edventure.com/PC2000**/.

The event is filling up fast, so get your registration in early!

RELEASE 1.0

SUBSCRIPTION FORM

☐ Please enter my subscription to *Release 1.0* at the rate of $795 per year in the U.S. and Canada. Overseas subscriptions are $850, airmail postage included. Payment must be enclosed. Satisfaction guaranteed or your money back.

☐ Please update my address as listed below.

Name
Title
Company
Address
City State Zip
Country

Telephone Fax
E-mail URL

☐ Check enclosed
☐ Charge my
  ☐ American Express ☐ Master Card ☐ Visa
  Card Number Expiration Date
  Name and Billing Address
  Signature

☐ Please send me information on your multiple copy rate.

Please fill in the information above and send to:

EDventure Holdings Inc.
104 Fifth Avenue, 20th Floor
New York, NY 10011

If you have any questions, please contact us at (212) 924-8800; fax (212) 924-0240; e-mail us@edventure.com; www.edventure.com.

Daphne Kis
Publisher

12-99